

Control What You Include!

Server-Side Protection against Third Party Web Tracking

Dolière Francis Somé, Nataliia Bielova, and Tamara Rezk

Université Côte d’Azur, Inria
`{doliere.some,tamara.rezk,nataliia.bielova}@inria.fr`

Abstract. Third party tracking is the practice by which third parties recognize users accross different websites as they browse the web. Recent studies show that 90% of websites contain third party content that is tracking its users across the web. Website developers often need to include third party content in order to provide basic functionality. However, when a developer includes a third party content, she cannot know whether the third party contains tracking mechanisms. If a website developer wants to protect her users from being tracked, the only solution is to exclude any third-party content, thus trading functionality for privacy. We describe and implement a privacy-preserving web architecture that gives website developers a control over third party tracking: developers are able to include functionally useful third party content, the same time ensuring that the end users are not tracked by the third parties.

1 Introduction

Third party tracking is the practice by which third parties recognize users accross different websites as they browser the web. In recent years, tracking technologies have been extensively studied and measured [25, 28, 31, 36, 24, 33] – researchers have found that third parties embedded in websites use numerous technologies, such as third-party cookies, HTML5 local storage, browser cache and device fingerprinting that allow the third party to recognize users across websites [37] and build browsing history profiles. Researchers found that more than 90% of Alexa top 500 websites [36] contain third party web tracking content, while some sites include as much as 34 distinct third party contents [30].

But why do website developers include so many third party contents (that may track their users)? Though some third party content, such as images and CSS [3] files can be copied to the main (first-party) site, such an approach has a number of disadvantages for other kinds of content. *Advertisement* is the base of the economic model in the web – without advertisements many website providers will not be able to financially support their website maintenance. *Third party JavaScript libraries* offer extra functionality: though copies of such libraries can be stored on the main first party site, this solution will sacrifice maintenance of these libraries when new versions are released. The developer would need to manually check the new versions. *Web mashups*, as for example applications that

use hotel searching together with maps, are actually based on reusing third-party content, as well as maps, and would not be able to provide their basic functionality without including the third-party content.

Including JavaScript libraries, content for mashups or advertisements means that the web developers cannot provide to the users the guarantee of non-tracking. Hence, the promise to provide privacy has a very high cost because there are no existing automatic tools to maintain control of third party tracking on the website. To keep a promise of non-tracking, the only solution today is to exclude any third-party content¹, thus trading functionality for privacy.

In this paper, we present a new Web application architecture that allows web developers to gain control over certain types of third party content. Our solution is based on the automatic rewriting of the web application in such a way that the third party requests are redirected to a trusted web server, with a different domain than the main site. This trusted web server may be either controlled by a trusted party, or by a main site owner – it is enough that the trusted web server has a different domain. A trusted server is needed so that the user’s browser will treat all redirected requests as third party requests, like in the original web application. The trusted server automatically eliminates third-party tracking cookies and other technologies.

In summary our contributions are:

- A classification of third party contents that can and cannot be controlled by the website developer.
- An analysis of third party tracking capabilities – we analyse two mechanisms: recognition of a web user, and identification of the website she is visiting ².
- A new architecture that allows to include third party content in web applications and eliminate stateful tracking.
- An implementation of our architecture, demonstrating its effectiveness at preventing stateful third party tracking in several websites.

2 Background and Motivation

Third party web tracking is the ability of a third party to re-identify users as they browse the web and record their browsing history [31]. Tracking is often done with the purpose of web analytics, targeted advertisement, or other forms of personalization. The more a third party is prevalent among the websites a user interacts with, the more precise is the browsing history collected by the tracker. Tracking has often been conceived as the ability of a third party to recognize the web user. However, for successful tracking, each user request should contain two components:

User recognition is the information that allows tracker to recognize the user; **Website identification** is the website which the user is visiting.

¹ For example, see <https://duckduckgo.com/>. ² Tracking is often defined as the ability of a third party to recognize a user through different websites. However, being able to identify the websites a user is interacting with is equally crucial for the effectiveness of tracking.




Fig. 1. Third Party Tracking

For example, when a user visits `news.com`, the browser may make additional requests to `facebook.com`, as a result, Facebook learns about the user's visit to `news.com`. Figure 1 shows a hypothetical example of such tracking where `facebook.com` is the third party.

Consider that a third party server, such as `facebook.com` hosts different contents, and some of them are useful for the website developers. The web developer of another website, say `mysite.com`, would like to include such *functional* content from Facebook, such as Facebook "Like" button, an image, or a useful JavaScript library, but the developer does not want its users to be tracked by Facebook. If the web developer simply includes third party Facebook content in his application, all its users are likely to be tracked by cookie-based tracking. Notice that each request to `facebook.com` also contains an HTTP Referrer header, automatically attached by the browser. This header contains the website URL that the user is visiting, which allows Facebook to build user's browsing history profile.

The example demonstrates cookie-based tracking, which is extremely common [36]. Other types of third party tracking, that use other client-side storage mechanisms, such as HTML5 LocalStorage, or cache, and device fingerprinting that do not require any storage capabilities, are also becoming more popular [25].

Web developer perspective A web developer may include third party content in her webpages, either because this content *intentionally* tracks users (for example, for targeted advertising), or because this content is important for the functioning of the web application. We therefore distinguish two kinds of third party contents from a web developer perspective: *tracking* and *functional*. *Tracking* content is *intentionally* embedded by website owner for tracking purposes. *Functional* content is embedded in a webpage for other purposes than tracking: for example, JavaScript libraries that provide additional functionality, such as jQuery, or other components, such as maps. In this work, we focus on *functional* content and investigate the following questions:

- What kind of third content is possible to control from a server-side (web developer) perspective?
- How to eliminate the two components of tracking (user recognition and website identification) from the functional third party that the website embeds?

2.1 Browsing Context

Browsers implement different specifications to securely fetch and aggregate third party content. One widely used approach is the the *Same Origin Policy (SOP)* [14], a security mechanism designed for developers to isolate legacy content from potentially untrusted third party content. An origin is defined as scheme, host and port number, of the URL³ of the third party content.

When a browser renders a webpage delivered by a first party, the page is placed within a *browsing context* [2]. A browsing context represents an instance of the browser in which a document such as a webpage is displayed to a user, for instance browser tabs, and popup windows. Each browsing context contains 1) a copy of the browser properties (such as browser name, version, device screen etc), stored in a specific object; 2) other objects that depend on the origin of the document according to SOP. For instance, the object `document.cookie` gives the cookies related to the origin of the current context.

In-context and cross-context content Certain types of content embedded in a webpage, such as images, links, and scripts, are associated with the context of the webpage, and we call them *in-context* content. Other types of content, such as `<iframe>`, `<embed>`, and `<object>` tags are associated with their own browsing context, and we call them *cross-context* content. Usually, cross-context content, such as `<iframe>` elements, cannot be visually distinguished from the webpage in which they are embedded, however they are as autonomous as other browsing contexts, such as tabs or windows. Table 1 shows different third party contents and their execution contexts.

	HTML Tags	Third Party Content
in-context	<code><link></code>	Stylesheets
	<code></code>	Images
	<code><audio></code>	Audios
	<code><video></code>	Videos
	<code><form></code>	Forms
	<code><script></code>	Scripts
cross-context	<code><(i)frame></code> , <code><frameset></code> , <code><a></code>	Web pages
	<code><object></code> , <code><embed></code> , <code><applet></code>	Plugins and Web pages

Table 1. Third party content and execution context.

The Same Origin Policy manages the interactions between different browsing contexts. In particular, it prevents in-context scripts from interacting with the content from a cross-context content in case their origins are different. To communicate, both contexts should rely on inter-frame communications APIs such as `postMessage` [12].

³ <https://www.w3.org/TR/url/>

2.2 Third Party Tracking

In this work, we consider only stateful tracking technologies – they require an identifier be stored client-side, the most common storage mechanism is cookies, but others, such as HTML5 LocalStorage and browser cache are also stateful tracking mechanisms. Figure 2 presents the well-known stateful tracking mechanisms. We distinguish two components necessary for successful tracking: user recognition and website identification. For each component, we describe the capabilities of in-context and cross-context. We also distinguish *passive tracking* (done through HTTP headers) and *active tracking* (through JavaScript or plugin script execution).

		User Recognition		Website Identification	
		Passive	Active	Passive	Active
<i>in-context</i>					
	HTTP cookies Cache-Control Etag Last-Modified		-	Referer Origin	<code>document.URL</code> <code>document.location</code> <code>window.location</code>
<i>cross-context</i>			Flash LSOs <code>document.cookie</code> <code>window.localStorage</code> <code>window.indexedDB</code>	Referer	<code>document.referrer</code>

Fig. 2. Stateful tracking mechanisms

In-context tracking In-context third party content is associated with the browsing context of the webpage that embeds it (see Table 1).

Passively, such content may use HTTP header to recognize the user and identify the visited website. When a webpage is rendered, the browser sends a request to fetch all third party contents embedded in the page. The response from the third party with the requested content may contain HTTP headers that may be used for tracking. For example, *Set-cookie* HTTP header tells the browser to save the third party cookies, that will be later automatically attached to every request to this third party in the Cookie header. Etag HTTP header and other cache mechanisms like *Last-Modified* and *Cache-Control* HTTP headers may also be used to store user identifier [37]. To identify the visited website, a third party can either check the *Referer* HTTP header, automatically attached by the browser, or an *Origin* header⁴.

Actively, in-context third party content cannot use browser storage mechanisms, such as cookies or HTML5 Local Storage associated to the third party

⁴ Origin header is also automatically generated by the browser when the third party content is trying to access data using Cross-Origin Resource Sharing [4] mechanism.

because of the limitations imposed by the SOP (see Section 2.1). For example, if a third party script uses `document.cookie` API, it is able only to read the cookies of the main website, but not those associated to the third party. This allows tracking within the main website but does not allow tracking cross-sites [36]. For website identification, third party active content, such as JavaScript, can use several APIs, such as `document.location` and others.

Cross-context tracking Cross-context third party content, such as `iframe`, is associated with the browsing context of the third party that provided this content.

Passively, the browser may transmit HTTP headers used for user recognition and website identification, just like with the in-context third party content. Every third-party request for cross-context content will contain the URL of the embedding webpage in its *Referer* header. Note that this is true only for the cross-context content, say an `<iframe>`, directly embedded in the webpage. Within the `iframe`, there may be additional third party contents. Since they are not embedded directly in the webpage, and because the `iframe` is an autonomous though nested browsing context, requests to fetch contents embedded within this context will carry, not the URL of the webpage, but that of the `iframe` in their *Referer* header, and the origin of the `iframe` in their CORS requests *Origin* header.

Actively, cross-context third party content can use a number of APIs to store user identifier in the browser. These APIs include cookies (`document.cookie`), HTML5 LocalStorage (`document.localStorage`), IndexedDB, and Flash Local Stored Objects (LSOs). For website identification, `document.referrer` API can be used – it returns the value of HTTP Referrer header transmitted to the third party when the third party content was fetched. Because cross-context third party is associated with its own browsing context, it is able to embed even more third party contents within this cross-context.

Combining in-context and cross-context tracking Imagine a third party script from `third.com` embedded in a webpage – according to the context and to the SOP, it is in-context. If the same webpage embeds another third party content from `third.com`, which is cross-context, then because of SOP, such script and `iframe` cannot interact directly. However, script and `iframe` can still communicate through inter-frame communication APIs such `postMessage` [12].

This communication between different contexts allow them to exchange the user identifiers and the website that the user visits. Efficient implementation of such combination of tracking may profit from easily implementable user recognition by cross-context code using, say `document.cookie`, and website identification by in-context through various APIs such as `document.location`. For example, social widgets, such as Facebook "Like" button, or Google "+1" button, may be included in the webpages as a script. When the social widget script is executed on the client-side, it loads additional scripts, and new browsing contexts (`iframes`) allowing the third party to benefit from both in-context and cross-context capabilities to track users.

3 Privacy-preserving Web Architecture

For third party tracking to be effective, it is necessary that it has two capabilities: 1) it is able to identify the website in which it is embedded, and 2) to recognize the user interacting with that website. Disabling only one of these two capabilities for a given third party already prevents tracking. In order to mitigate the stateful tracking (see Section 2), we make the following design choices in our architecture:

1. **In-context content: prevent only user recognition.** Preventing passive user recognition for in-context content, such as images, forms and scripts is possible by removing HTTP headers such as *Set-cookie*, *ETag* and others. However, it is particularly difficult to remove active website identification because trying to alter or redefine `document.location` and `window.location` APIs, will cause the main page to reload.
2. **Cross-context content: prevent only website identification.** We prevent passive website identification by instructing the browser not to send HTTP *Referer* header along with requests to fetch a cross-context content. Therefore, when the cross-context gets loaded, active website identification is impossible. Indeed, executing `document.referrer` returns not the URL of the embedding page, but an empty string. Because of the limitations of the SOP, a website owner has no control over the cross-context third party content, such as iframes. Therefore, it is not possible to modify the results of storage access APIs, such as `document.cookie`. We discuss other possibilities to block such APIs in Section 4.3.
3. **Prevent communication between in-context and cross-context contents.** Our architecture proposes a way to block such communications that can be done by `postMessage` API. We discuss the limitations of this approach in Section 4.3.

To help web developers keep their promises of non-tracking and still include third-party content in their web applications, we propose a new Web application architecture. This architecture has the capability to 1) automatically rewrite all the third party *in-context* content of a Web application, 2) redirect the third party HTTP requests issued by the *in-context* content, and 3) remove/disable known *stateful* tracking mechanisms (see Section 2) for such third party content and requests. 4) It also rewrites and redirects cross-context requests so as to prevent website identification and communication with in-context scripts.

Figure 3 provides an overview of our web application architecture, that introduces two new components that are fully controlled by the website owner:

Rewrite Server (Section 3.1) acts like a reverse proxy⁵ for the original web server. It rewrites the web pages in such a way that all the third party requests are redirected through the Middle Party Server before reaching the intended third party server.

⁵ https://en.wikipedia.org/wiki/Reverse_proxy

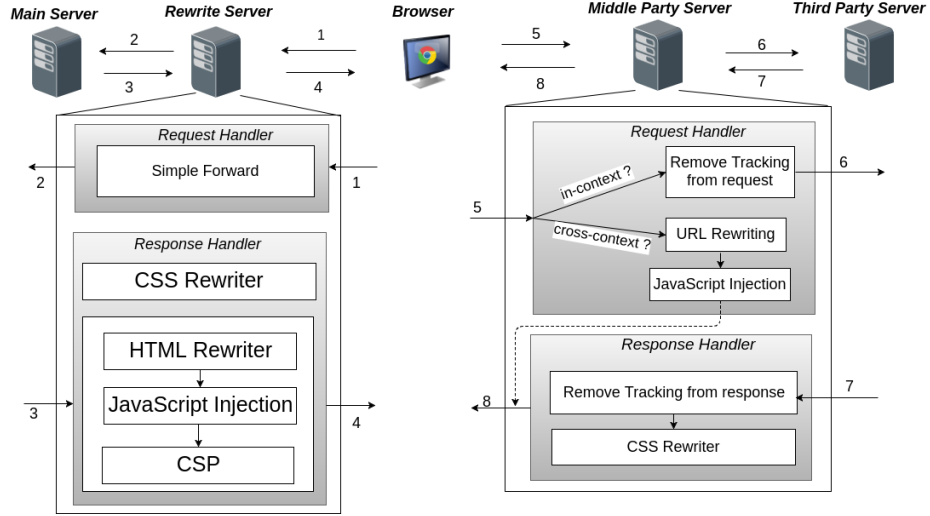


Fig. 3. Privacy-Preserving Web Architecture


Middle Party Server (Section 3.1) is at the core of our solution since it intercepts all browser third party requests, removes tracking, then forwards them to the intended third parties. When they reply, it also removes tracking information and forwards the responses back to the browser. On one hand, it hides the third party destination from the browser, and therefore prevents the browser from attaching third party HTTP cookies to such requests. Because the browser will still attach some tracking information to the requests, such as *ETag*, and *Referer* headers, Middle Party Server will also remove this information when forwarding the requests to the third party. This prevents passive user identification for in-context third party contents.

On the other hand, the Middle Party Server prevents website identification for cross-context contents and communication with in-context scripts. This is done by placing the cross-context within another cross-context controlled by the Middle Party server as illustrated by Figure 4. For instance, if an iframe was to be embedded within a webpage, it is placed within another iframe that belongs to the Middle Party. The Middle Party then instructs the browser not to send *Referer* header while loading the iframe, which prevents passive and active website identification once it is loaded. Since the iframe is nested within a iframe that belongs to Middle Party, this hides its reference to in-context scripts (see Figure 4). Therefore, it is prevented from communicating with in-context scripts in the main webpage.

3.1 Rewrite Server

The goal of the Rewrite Server is to rewrite the original content of the requested webpages in such a way that all third party requests will be redirected to the

API	Content
<code>document.createElement</code>	inject contents from Table 1
<code>document.write</code>	any content
<code>window.open</code>	Web pages(popups)
<i>Image</i>	images
<i>XMLHttpRequest</i>	any data
<i>Fetch, Request</i>	any content
<i>Event Source</i>	stream data
<i>WebSocket</i>	websocket data

Table 2. Embedding Dynamic Third Party Contents**Fig. 4.** Prevent Combining in-context and cross-context tracking

Middle Party Server. It consists of three main components: static HTML rewriter for HTML pages, static CSS rewriter and JavaScript injection component. Into each webpage, we inject a JavaScript code that insures that all the dynamically generated third party content is redirected to the Middle Party Server.

HTML and CSS Rewriter rewrites the URLs of static third party contents embedded in original web pages and CSS files in order to redirect them to the Middle Party Server. For example, the URL of a third-party script source `http://third.com/script.js` is written so that it is instead fetched through the Middle Party Server: `http://middle.com/?src=http://third.com/script.js`.

JavaScript Injection. The Rewrite Server also injects a script in an original webpage, that controls APIs used to inject dynamic contents. This injected script rewrites third party contents which are dynamically injected in webpages after they are rendered on the client-side. Table 2 shows APIs that can be used to dynamically inject third party content within a webpage that we control using the injected script.

A **Content Security Policy (CSP)** [41] is injected in the response header for each webpage in order to prevent third parties from bypassing the rewriting and redirection to the Middle Party Server. A CSP delivered with the webpage controls the resources of the page. It allows to specify which resources are allowed to be loaded and executed in the page. By limiting the resource origins to only those from the Middle Party Server and the website own domain, we prevent third parties from bypassing our redirection to the Middle Party Server.

3.2 Middle Party

The main goal of the Middle Party is to proxy the requests and responses between browsers and third parties in order to remove tracking information exchanged

between them. For in-context contents, it removes any user recognition as well as website identification information. For cross-context contents, it takes care of preventing website identification and communication with in-context scripts.

In-Context Contents are scripts, images, etc. (see Table 1). Since a third party script from `http://third.com/script.js` is rewritten by the Rewrite Server to `http://middle.com/?src=http://third.com/script.js`, it is fetched through the Middle Party Server. When the middle party receives such a request URL from the browser, it takes the following steps. **Remove Tracking from request** that are set by the browser as HTTP headers. Among those headers are *Cookie*, *Etag*, *If-Modified-Since*, *Cache-Control*, *Referer*. Next, it makes a request to the third party in order to get the content of the script `http://third.com/script.js`. **Remove Tracking from response** returned by the third party. The headers that the third party may send are *Set-Cookie*, *Etag*, *Last-Modified*, *Cache-Control*. **CSS Rewriter** rewrites the response if the content is a CSS file. Finally, the response is returned back to the browser.

Cross-context contents are iframes, links, popups, etc. (see Table 1). For instance, a third party iframe from `http://third.com/page.html` is rewritten to `http://middle.com/?emb=http://third.com/page.html`. When the Middle Party Server receives such a request URL from the browser, it takes the following actions: **URL Rewriting**: instead of fetching directly the content of `http://third.com/page.html`, the Middle Party Server generates a content in which it puts the URL of the third party content as a hyperlink. ``. The most important part of this content is in the `rel` attribute value. Therefore, `noreferrer noopener` instructs the browser not to send the *Referer* header when the link `http://third.com/page.html` is followed client-side. **JavaScript injection** module adds a script to the content so that the link gets automatically followed once the response is rendered by the browser. Once the link is followed, the browser fetches the third party content directly on the third party server, without going through the Middle Party server anymore. Nonetheless, it does not include the *Referer* header for identifying the website. Therefore, the `document.referrer` API also returns an empty string inside the iframe context. This prevents it from identifying the website.

The third party server response is placed within a new iframe nested within a context that belongs to the Middle Party, and not directly within the site webpage. This prevents in-context scripts and the cross-context contents from exchanging tracking information as illustrated by Figure 4.

4 Implementation

We have implemented both the Rewrite Server and the Middle Party Server as full Node.js [10] web servers supporting *HTTP(S)* protocols and web sockets. Implementation details are available at <https://webstats.inria.fr/sstp/>.

4.1 Rewrite Server

Simple Forward: requests that arrive to the Rewrite server are simply forwarded to the main server.

HTML Rewriter is implemented with Jsdom HTML parser [8] and **CSS Rewriter** using a CSS parser [5] for Node.js. **JavaScript injection** is done at the end of rewriting webpages. The code script injected is available at <https://webstats.inria.fr/sstp/dynamic.js>. **CSP** set on webpages only whitelists the website own domain and the Middle Party. It also prevents third party plugins.

```
1 | Content-Security-Policy: default-src 'self' 'middle.com';
   | object-src 'self';
```

4.2 Middle Party

In-Contexts Contents. Remove Tracking from requests component removes tracking information from in-context third party requests (See Section 3). The requests are then forwarded to the original third party server, to fetch the third party content. **Remove Tracking from responses** : Tracking information that are set by third parties in the responses, are removed. See Section 3 for details about information that are removed. **CSS Rewriter**: as in the case of the Rewrite Server, this component is implemented using a CSS parser [5] for Node.js for rewriting CSS files.

Cross-Context Contents. URL Rewriting If the cross-context URL was <http://third.com/page.html>, this URL is rewritten to

```
1 | <a href="http://third.com/page.html" rel="noreferrer
   | noopener" target=""></a>.
```

JavaScript injection : the content injected is as followed.

```
1 | var third_party = document.getElementsByTagName("a")[0];
2 | if(window.top == window.self){
3 |     third_party.target = "_blank";
4 |     third_party.click();
5 |     window.close();
6 | }else{
7 |     var iframe = document.createElement("iframe");
8 |     iframe.name = "iframe_target";
9 |     document.body.appendChild(iframe);
10 |    third_party.target = "iframe_target";
11 |    third_party.click();
12 | }
```

Both the rewritten URL and the injected script are returned as a response to the browser which renders it. The injected script gets executed within a context that belongs to the Middle Party. If the original cross-context third party content was to be loaded inside an iframe, the injected script creates an iframe in which the

original third party content is loaded. However, if it was to be opened inside a new tab, the injected script opens a new tab in which the third party content is loaded. In both cases, while the cross-context content is loaded, the browser does not send the *Referer* header. This makes the value of `document.referrer` empty inside the cross-context preventing it from identifying the website. Finally, since those cross-context are loaded by the injected script from a context that belongs to the Middle Party, in-context scripts cannot communicate with the cross-context contents to exchange tracking information.

4.3 Discussion and Limitations

Our approach suffers from the following limitations. First, our implementation prevents cross-context and in-context contents from communicating with each other using `postMessage` API. However, in-context third party script can identify the website a user visits via `document.location.href` API. Then the script can include the website URL, say `http://main.com`, as a parameter of the URL of a third party iframe, for example `http://third.com/page.html?ref=http://main.com` and dynamically embed it in the webpage. In our architecture, this URL is rewritten and routed to the Middle Party. Since, the Middle Party Server does not inspect URL parameters, this information will reach the third party even though the *Referer* is not sent with cross-context requests.

Another limitation is that of dynamic CSS changes. For instance, changing the background image style of an element in the webpage is not captured by the dynamic rewriting script injected in webpages. Therefore, if the image was a third party image, the CSP will prevent it from loading.

Performance overhead There is a performance cost associated with the Rewrite Server. Rewriting contents server-side and browser-side is also expensive in terms of performance. Middle Party Server may also lead to performance overhead especially for webpages with numerous third party contents. We believe that server-side caching mechanisms may help to speed up responsiveness.

Extension to stateless tracking Even though this work did not address stateless tracking, such as device fingerprinting, our architecture already hides several fingerprintable device properties and can be extended to several others: 1) The redirection to the Middle Party anonymizes the real IP addresses of users; 2) Some stateless tracking APIs such as `window.navigator`, `window.screen`, and `HTMLCanvasElement` can be easily removed or randomized from the context of the webpage to mitigate in-context fingerprinting.

Possibility to blocking active user recognition in cross-context With the prevalence of third party tracking on the web, we have shown the challenges that a developer will face towards mitigating that. The sandbox attribute for iframes help prevent access to security-sensitive APIs. As tracking has become a hot concern, we suggest that similar mechanisms can help first party websites tackle third party tracking. The sandbox attribute can for instance be extended with specific values to tackle tracking. Nonetheless, the sandbox attribute can be used to prevent cross-context from some stateful tracking mechanisms [9].

5 Evaluation and Case Study

Demo website We have set up a demo website that embeds a collection of third party contents, both in-context and cross-context. In-context contents include images, HTML5 audio and video, and a Google Map, which further loads dynamic contents such as images, fonts, scripts, and CSS files; a Youtube video as a cross-context content. Our demo website is accessible at <http://sstp-rewriteproxy.inria.fr>. When we deployed the Rewrite Server on <http://sstp-rewriteproxy.inria.fr>, the original server has been moved to <http://sstp-rewriteproxy.inria.fr:8080>, so that it is no longer directly accessible to users. The Middle Party server runs at <http://sstp-middleparty.inria.fr>.

Originally, when all the third parties were simply included in the main webpage, they may have also been tracking the website users (see Figure 1). After the deployment of our solution, we have been able to redirect all in-context third party contents to the Middle Party. We have been able to prevent the website identification in the cross-context Youtube video. In the Appendix, we show a screenshot of requests redirection to the Middle Party Server.

Real websites Since we did not have access to a real websites, we cannot install a Rewrite Server and to evaluate our solution. We therefore implemented a browser proxy based on a Node.js proxy [11], and included all the logic of the Rewrite Server within the proxy. The proxy is running at <http://sstp-rewriteproxy.inria.fr:5555>.

We then evaluated the solution on different kinds of websites: a news website <http://www.bbc.com>, an entertainment website <http://www.imdb.com>, and a shopping website <http://verbaudet.fr>. All three websites load content from various third party domains. In all websites, we rewrote all third party contents through the proxy (acts as Rewrite Server) and the Middle Party Server removed tracking information. Visually, we did not notice any change in the behaviors of the websites. We also interacted with them in a standard way (clicking on links on a news website, choosing products and putting them in the basket on the shopping website) and all the main functionalities of the websites was preserved.

Overall, these evaluation scenarios have helped us improve the solution, especially rewriting dynamically injected third party content. We believe that this implementation will even get better in the future when we convince to deploy it for some real websites.

6 Related Work

Many studies have demonstrated that third party tracking is very prevalent on the web today as well as the underlying tracking technologies [25, 36, 31, 28]. Lerner et al. [30] dusted the story of this practice for a period of twenty years. Trackers have been categorized according to either their business relationships with websites [31], their prominence [28, 25] or the user browsing profile that they can build [36]. Mayer and Mitchell [31] grouped tracking mechanisms in two

categories called statefull (cookie-based and super-cookies) and stateless (fingerprinting). It is rather intuitive to convince ourselves about the effectiveness of a statefull tracking, since the latter is based on unique identifiers that are set in users browsers. Nonetheless, the efficacy of stateless mechanisms has been extensively demonstrated. Since the pioneer work of Eckersley [24], new fingerprinting methods have been revealed in the literature [38, 22, 25, 19, 21, 17, 39, 33, 18]. A classification of fingerprinting techniques is provided in [40]. Those studies have contributed to raising public awareness of tracking privacy threats. Mayer and Mitchell [31] have shown that users are very sensitive to their online privacy, thus hostile to third party tracking. Englehardt et al. [26] have demonstrated that tracking can be used for surveillance purposes. The success of anti-tracking defenses is yet another illustration of users concern regarding tracking [32].

There are many defenses that try to protect users against third party tracking. First, major browser vendors do natively provide mechanisms for users to block third party cookies, browse in private mode. More and more privacy-browsers even take a step further, putting privacy as a design and implementation principle. Examples of such browsers are the Tor Browser [16], TrackingFree Browser [34] or Blink [29]. But the most popular defenses are by far browser extensions. Being tightly integrated to browsers, they provide additional privacy features that are not natively implemented in browsers. Well known extensions for privacy are Disconnect [6], Ghostery [7], Adblock [1], ShareMeNot [36], which is now part of PrivacyBadger [13], MyTrackingChoices [20], MyAdChoices [35]. Merzdovnik et al. [32] provide a large-scale study of anti-tracking defenses. Well known trackers such as advertisers, which businesses hugely depend on tracking, have also been taking steps towards limiting their tracking capabilities [31]. The W3C is pushing forward the Do Not Tracking standard [23, 27] for users to easily express their tracking preferences so that trackers may comply with them. To the best of our knowledge, we are the first to investigate how a website owner can embed third party content while preventing them from accidentally tracking users. The idea of proxying requests within a webpage is inspired by web service workers API [15], though the latter is still a working draft which is being tested in Mozilla Firefox and Google Chrome.

7 Conclusions

Most of the previous research analysed third party tracking mechanisms, and how to block tracking from a user perspective. In this work, we classified third party tracking capabilities from a website developer perspective. We proposed a new architecture for website developers that allows to embed third party contents while preserving users privacy. We implemented our solution, and evaluated it on real websites to mitigate stateful tracking.

References

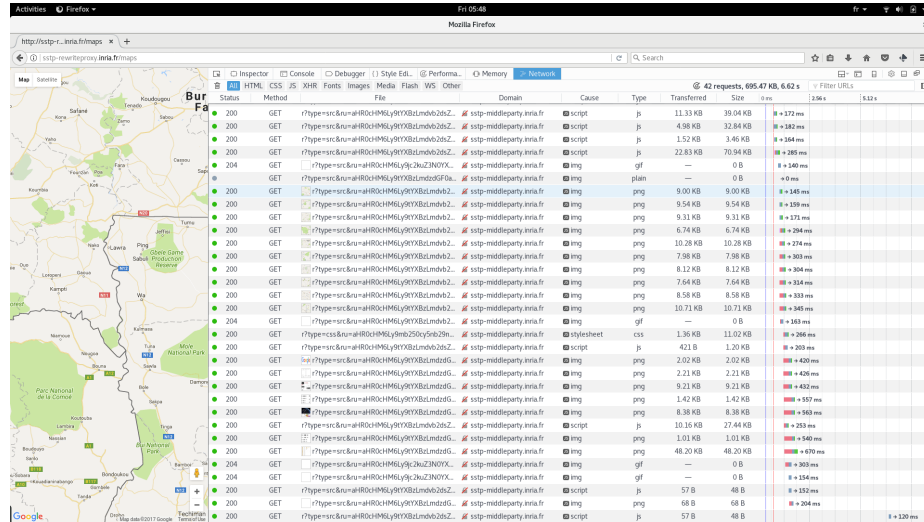
- [1] Adblock - Block Ads - Browse Safe. <https://getadblock.com/>.
- [2] Browsing Contexts. <https://www.w3.org/TR/html51/browsers.html>.
- [3] Cascading Style Sheets. <https://www.w3.org/Style/CSS/>.
- [4] Cross-origin-resource sharing. https://developer.mozilla.org/en-US/docs/Web/HTTP/Access_control_CORS.
- [5] CSS Parser for Node.js. <https://github.com/reworkcss/css>.
- [6] Disconnect. <https://disconnect.me/>.
- [7] Ghostery. <https://www.ghostery.com/>.
- [8] HTML Parser for Node.js. <https://github.com/tmpvar/jsdom>.
- [9] Iframe Sandbox Attribute. <https://www.w3.org/TR/2011/WD-html5-20110525/the-iframe-element.html#attr-iframe-sandbox>.
- [10] Node.js. <https://nodejs.org/en/>.
- [11] Node.js Proxy. <https://newspaint.wordpress.com/2012/11/05/node-js-http-and-https-proxy>.
- [12] PostMessage - Cross-Origin Iframe Secure Communication. <https://developer.mozilla.org/en-US/docs/Web/API/Window/postMessage>.
- [13] Privacy Badger - Electronic Frontier Foundation. <https://www.eff.org/fr/privacybadger>.
- [14] Same Origin Policy. https://www.w3.org/Security/wiki/Same_Origin_Policy.
- [15] Service Worker API. https://developer.mozilla.org/en-US/docs/Web/API/Service_Worker_API.
- [16] The Design and Implementation of the Tor Browser [Draft]. <https://www.torproject.org/projects/torbrowser/design/>.
- [17] E. Abgrall, Y. L. Traon, M. Monperrus, S. Gombault, M. Heiderich, and A. Ribault. XSS-FP: browser fingerprinting using HTML parser quirks. *CoRR*, abs/1211.4812, 2012.
- [18] G. Acar, C. Eubank, S. Englehardt, M. Juárez, A. Narayanan, and C. Díaz. The web never forgets: Persistent tracking mechanisms in the wild. In G. Ahn, M. Yung, and N. Li, editors, *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 674–689. ACM, 2014.
- [19] G. Acar, M. Juárez, N. Nikiforakis, C. Díaz, S. F. Gürses, F. Piessens, and B. Preneel. Fpdetective: dusting the web for fingerprinters. In A. Sadeghi, V. D. Gligor, and M. Yung, editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 1129–1140. ACM, 2013.
- [20] J. P. Achara, J. Parra-Arnau, and C. Castelluccia. Mytrackingchoices: Pacifying the ad-block war by enforcing user privacy preferences. *CoRR*, abs/1604.04495, 2016.
- [21] K. Boda, Á. M. Földes, G. G. Gulyás, and S. Imre. User tracking on the web via cross-browser fingerprinting. In P. Laud, editor, *Information Security Technology for Applications - 16th Nordic Conference on Secure IT Systems, NordSec 2011, Tallinn, Estonia, October 26-28, 2011, Revised Selected Papers*, volume 7161 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2011.
- [22] Y. Cao, S. Li, and E. Wijmans. (cross-)browser fingerprinting via os and hardware level features. In *24th Annual Network and Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, 26 February - 1 March, 2017*, 2017. To Appear.

- [23] N. Doty. Tracking Compliance and Scope, 2016. <https://www.w3.org/TR/tracking-compliance/>.
- [24] P. Eckersley. How unique is your web browser? In M. J. Atallah and N. J. Hopper, editors, *Privacy Enhancing Technologies, 10th International Symposium, PETS 2010, Berlin, Germany, July 21-23, 2010. Proceedings*, volume 6205 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2010.
- [25] S. Englehardt and A. Narayanan. Online tracking: A 1-million-site measurement and analysis. In E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 1388–1401. ACM, 2016.
- [26] S. Englehardt, D. Reisman, C. Eubank, P. Zimmerman, J. Mayer, A. Narayanan, and E. W. Felten. Cookies that give you away: The surveillance implications of web tracking. In A. Gangemi, S. Leonardi, and A. Panconesi, editors, *Proceedings of the 24th International Conference on World Wide Web, WWW 2015, Florence, Italy, May 18-22, 2015*, pages 289–299. ACM, 2015.
- [27] R. T. Fielding. Tracking Preference Expression (DNT), 2015. <https://www.w3.org/TR/tracking-dnt/>.
- [28] B. Krishnamurthy and C. E. Wills. Privacy diffusion on the web: a longitudinal perspective. In J. Quemada, G. León, Y. S. Maarek, and W. Nejdl, editors, *Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, April 20-24, 2009*, pages 541–550. ACM, 2009.
- [29] P. Laperdrix, W. Rudametkin, and B. Baudry. Beauty and the beast: Diverting modern web browsers to build unique browser fingerprints. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016*, pages 878–894. IEEE Computer Society, 2016.
- [30] A. Lerner, A. K. Simpson, T. Kohno, and F. Roesner. Internet jones and the raiders of the lost trackers: An archaeological study of web tracking from 1996 to 2016. In *25th USENIX Security Symposium (USENIX Security 16)*, Austin, TX, 2016. USENIX Association.
- [31] J. R. Mayer and J. C. Mitchell. Third-party web tracking: Policy and technology. In *IEEE Symposium on Security and Privacy, SP 2012, 21-23 May 2012, San Francisco, California, USA*, pages 413–427. IEEE Computer Society, 2012.
- [32] G. Merzdovnik, M. Huber, D. Buhov, N. Nikiforakis, S. Neuner, M. Schmiedecker, and E. Weippl. Block me if you can: A large-scale study of tracker-blocking tools. In *2nd IEEE European Symposium on Security and Privacy*, Paris, France, 2017. To appear.
- [33] N. Nikiforakis, A. Kapravelos, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna. Cookieless monster: Exploring the ecosystem of web-based device fingerprinting. In *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*, pages 541–555. IEEE Computer Society, 2013.
- [34] X. Pan, Y. Cao, and Y. Chen. I do not know what you visited last summer: Protecting users from stateful third-party web tracking with trackingfree browser. In *22nd Annual Network and Distributed System Security Symposium, NDSS 2015, San Diego, California, USA, February 8-11, 2015*. The Internet Society, 2015.
- [35] J. Parra-Arnau, J. P. Achara, and C. Castelluccia. Myadchoices: Bringing transparency and control to online advertising. *CoRR*, abs/1602.02046, 2016.
- [36] F. Roesner, T. Kohno, and D. Wetherall. Detecting and defending against third-party tracking on the web. In S. D. Gribble and D. Katabi, editors, *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation*,

- NSDI 2012, San Jose, CA, USA, April 25-27, 2012*, pages 155–168. USENIX Association, 2012.
- [37] A. Soltani, S. Canty, Q. Mayo, L. Thomas, and C. J. Hoofnagle. Flash cookies and privacy. In *AAAI spring symposium: intelligent information privacy management*, pages 158–163, 2010.
 - [38] O. Starov and N. Nikiforakis. Extended tracking powers: Measuring the privacy diffusion enabled by browser extensions. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3 - 7, 2017*, 2017. To Appear.
 - [39] N. Takei, T. Saito, K. Takasu, and T. Yamada. Web browser fingerprinting using only cascading style sheets. In L. Barolli, F. Khafa, M. R. Ogiela, and L. Ogiela, editors, *10th International Conference on Broadband and Wireless Computing, Communication and Applications, BWCCA 2015, Krakow, Poland, November 4-6, 2015*, pages 57–63. IEEE Computer Society, 2015.
 - [40] R. Upathilake, Y. Li, and A. Matrawy. A classification of web browser fingerprinting techniques. In M. Badra, A. Boukerche, and P. Urien, editors, *7th International Conference on New Technologies, Mobility and Security, NTMS 2015, Paris, France, July 27-29, 2015*, pages 1–5. IEEE, 2015.
 - [41] M. West, A. Barth, and D. Veditz. Content Security Policy Level 2. W3C Candidate Recommendation, 2015.

Appendix

Screenshot of the demo website map console.



Status	Method	File	Domain	Cause	Type	Size	Time	Size	Time
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	script	js	11.33 KB	39.04 KB	8 = 373 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	script	js	4.58 KB	32.84 KB	8 = 382 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	script	js	1.52 KB	3.46 KB	8 = 364 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	script	js	22.83 KB	70.94 KB	8 = 285 ms	
204	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	gif	0 B	0 B	8 = 140 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	9.00 KB	9.00 KB	8 = 145 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	9.54 KB	9.54 KB	8 = 139 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	9.31 KB	9.31 KB	8 = 171 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	6.74 KB	6.74 KB	8 = 294 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	10.20 KB	10.20 KB	8 = 274 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	7.88 KB	7.88 KB	8 = 303 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	8.12 KB	8.12 KB	8 = 304 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	7.64 KB	7.64 KB	8 = 314 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	8.58 KB	8.58 KB	8 = 333 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	10.71 KB	10.71 KB	8 = 345 ms	
204	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	gif	0 B	0 B	8 = 135 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	stylesheet	css	1.36 KB	11.02 KB	8 = 206 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	script	js	421 B	1.20 KB	8 = 203 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	2.02 KB	2.02 KB	8 = 420 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	2.21 KB	2.21 KB	8 = 426 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	9.21 KB	9.21 KB	8 = 432 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	1.42 KB	1.42 KB	8 = 557 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	8.38 KB	8.38 KB	8 = 540 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	script	js	10.16 KB	27.44 KB	8 = 253 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	1.01 KB	1.01 KB	8 = 540 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	48.20 KB	48.20 KB	8 = 470 ms	
204	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	gif	0 B	0 B	8 = 303 ms	
204	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	gif	0 B	0 B	8 = 154 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	script	js	57 B	48 B	8 = 152 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	img	png	68 B	68 B	8 = 204 ms	
200	GET	/r?type=s&u=aHROCHM&ly9tXB&Lmd&2&Z...	stsp-middleparty.inna.fr	script	js	57 B	48 B	8 = 120 ms	

Fig. 5. Screenshot of the Browser console